Spatiotemporal dynamics and reliable computations in recurrent spiking neural networks

Ryan Pyle[1] and Robert Rosenbaum[1,2]

[1]Department of Applied and Computational Mathematics and Statistics, University of Notre Dame, Notre Dame, IN 46556
[2]Interdisciplinary Center for Network Science and Applications, University of Notre Dame, Notre Dame, IN 46556

Randomly connected networks of excitatory and inhibitory spiking neurons provide a parsimonious model of neural variability, but are notoriously unreliable for performing computations. We show that this difficulty is overcome by incorporating the well-documented dependence of connection probability on distance. Spatially extended spiking networks exhibit symmetry-breaking bifurcations and generate spatiotemporal patterns that can be trained to perform dynamical computations.

arXiv:1611.01557v1 [q-bio.NC] 4 Nov 2016

Biological neuronal networks exhibit irregular and asynchronous activity [1, 2] that is often modeled using randomly connected networks of excitatory and inhibitory spiking neurons. In these models, an approximate balance between excitation and inhibition creates a push-pull dynamic that combines with random connectivity and chaotic network dynamics to produce asynchronous-irregular spiking activity similar to that observed in experimental recordings [3–6].

Despite their ability to explain the genesis of neural variability, asynchronous-irregular spiking network models have a critical shortcoming: Their microscopic dynamics – at the level of spike times – are intricate and nonlinear, but largely unreliable [3, 4, 7, 8]. Their macroscopic dynamics – at the level of firing rates – are reliable, but primarily track network input [3, 4, 6, 9, 10]. Biological neural networks generate reliable, intricate responses to simple sensory inputs, for example to produce motor output [11, 12]. This raises the question of how neural circuits reliably produce intricate firing rate dynamics for dynamical computations.

In this letter, we show that the limited dynamical complexity of firing rates in asynchronous-irregular spiking networks is overcome by incorporating the widely reported dependence of connection probability on distance [13–16]. Spiking networks with a spatial topology can undergo symmetry-breaking Turing-Hopf bifurcations [17, 18] to generate intricate spatiotemporal dynamics that can be trained to perform computations.

Results Following previous work [9], we consider a recurrent neural network with 4×10^4 excitatory (e) and 10^4 inhibitory (i) model neurons arranged uniformly on a square-shaped domain, $\Gamma = [0,1] \times [0,1]$, with periodic boundaries, *i.e.* a torus. The synaptic input current to neuron j in population $a = \mathrm{e}, \mathrm{i}$ is given by

$$I_j^a(t) = \sum_{k=1}^{N_\mathrm{e}} J_{jk}^{a\mathrm{e}} \sum_n \delta(t - t_{nk}^\mathrm{e}) + \sum_{k=1}^{N_\mathrm{i}} J_{jk}^{a\mathrm{i}} \sum_n \delta(t - t_{nk}^\mathrm{i}) + F_j^a(t)$$

where t_{nk}^b is the nth spike of neuron k in population $b = \mathrm{e}, \mathrm{i}$. Spikes are determined by a leaky integrate-and-fire dynamic [42]. To model the widely observed distance-dependence of connection probability [13–16], the synaptic weight from a neuron at coordinates $\mathbf{y} \in \Gamma$ in population b to a neuron at $\mathbf{x} \in \mathbf{\Gamma}$ in population a is chosen randomly according to

$$J_{jk}^{ab} = \begin{cases} j_{ab} & \text{with prob. } p_{ab}(\mathbf{x} - \mathbf{y}) \\ 0 & \text{otherwise} \end{cases},$$

where $p_{ab}(\mathbf{u}) = \overline{p}_{ab} G(\mathbf{u}; \sigma_b)$ and $G(\mathbf{u}; \sigma_b)$ is a two-dimensional wrapped Gaussian with width σ_b [9].

We first simulated a network in which external inputs were constant across space and time and inhibitory projections were more localized than excitatory projections (Fig. 1a). Even though the model is deterministic, spiking activity was irregular and asynchronous with no coherent spatial patterning (Fig. 1b, Supplementary Figures 1,2 and Supplementary Animation). This biologically realistic spike-timing variability is driven by chaos-like dynamics [3, 4, 7–9].

Despite the complexity of spike-timing dynamics, firing rates are amenable to mean-field analysis under a diffusion approximation [9]. The mean input, $\vec{\mu}(\mathbf{x}) = [\mu_\mathrm{e}(\mathbf{x})\ \mu_\mathrm{i}(\mathbf{x})]^T$, to e and i neurons near $\mathbf{x} \in \Gamma$ is

$$\vec{\mu}(\mathbf{x}) = \iint_\Gamma W(\mathbf{u}) \vec{r}(\mathbf{x} - \mathbf{u}) d\mathbf{u} + \vec{F}(\mathbf{x})$$

where $\vec{r}(\mathbf{x}) = [r_\mathrm{e}(\mathbf{x})\ r_\mathrm{i}(\mathbf{x})]^T$ is the average firing rate and $\vec{F}(\mathbf{x})$ the feedforward input to neurons near $\mathbf{x} \in \Gamma$. The matrix kernel $W(\mathbf{u})$ captures synaptic divergence and similarly for the input variance, $v(\mathbf{x}) = \iint U(\mathbf{u}) \vec{r}(\mathbf{x} - \mathbf{u}) d\mathbf{u}$ [43]. The mapping from input statistics to rates, $\vec{r} = \phi(\vec{\mu}, \vec{v})$, is computed using a Fokker-Planck formalism so fixed point rates can be computed numerically [19–22]. When $\vec{F}(\mathbf{x}) = \vec{F}$ is spatially uniform, so are fixed point rates [9].

To estimate local rates from simulations, we partitioned the network into 100 squares, then averaged and low-pass filtered the spike trains of the 400 excitatory neurons in each square (Fig. 1a). This yields a readout of the local instantaneous firing rates and also models synaptic output that the network would send to downstream neural populations. These local rate readouts closely matched the fixed point firing rates computed numerically from the diffusion approximation (Fig. 1c) and fluctuations in the rates were consistent with asynchronous, Poisson-like spike timing variability [44].

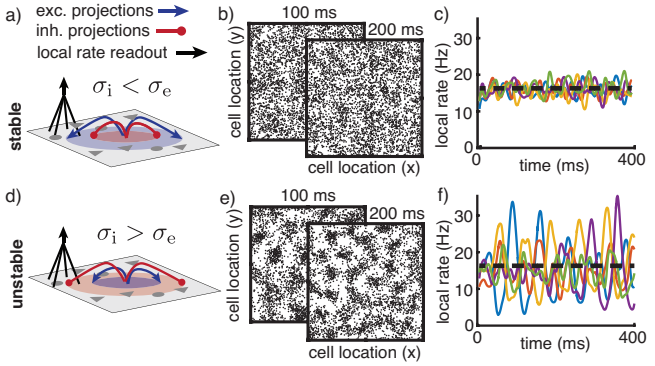

FIG. 1: **Intrinsic dynamics in a spatially extended spiking network.** **a)** Schematic of spatially extended spiking network model. Excitatory and inhibitory neurons arranged on a square project randomly to one another. Lateral excitatory projections (blue) are are longer on average than lateral inhibitory projections (red; $\sigma_e = 0.1$, $\sigma_i = 0.05$, $F^e(t) = 3$ V/s, $F^i(t) = 2.3$ V/s). Each local rate readout (black) is computed by averaging the activity of all excitatory neurons within a square patch then low-pass filtering with a Gaussian kernel ($\sigma = 5$ ms). **b)** Raster plot snapshots over 5 ms time windows starting at $t = 100$ and 200 ms. **c)** Five randomly chosen local rate readouts. Dashed black line shows numerically computed fixed point rate. **d-f)** Same as a-c, but inhibitory projections are broader than excitatory ($\sigma_e = 0.05$, $\sigma_i = 0.1$). See Supplemental Animation for animated raster plots.

Broad lateral inhibition is known to induce spatial pattern formation [9, 23–27]. We next modified the network so that inhibitory projections were broader than excitatory projections (Fig. 1d). This produced a dramatic change in the spiking activity, with spatially uniform activity giving way to intricate, asymmetric spatiotemporal activity patterns (Fig. 1e and Supplementary Animation), despite the spatial symmetry of connection probabilities in the network. These spatiotemporal patterns were reflected in the local rate readouts by irregular high-amplitude fluctuations (Fig. 1f). Despite their differences, both networks produced asynchronous, irregular spike trains with an approximate balance between excitation and inhibition [45].

The network with broad inhibition and the network with local inhibition share the same spatially uniform fixed point under the diffusion approximation, but rates strongly deviated from this fixed point when inhibition was broader. We therefore conjectured that the fixed point was stable for the simulation with local inhibition and unstable when inhibition was broader.

Spatially extended neural networks are often described using integro-differential equations of the form

$$\tau \partial \vec{r}/\partial t = -\vec{r} + \phi(\vec{\mu}, \vec{v}) \tag{1}$$

or of a similar form. As in previous work [9], this approach predicted stability of uniform firing rates for the network with broad inhibition, despite the patterns observed in simulations [46]. We conjectured that firing rate dynamics observed when inhibition was broad arose in part from a resonance in neurons' membrane and spiking dynamics [28] that is not captured by Eq. (1). To account for this resonance, we generalized the stability analysis from recent work [29] to spatial networks. Linear response theory gives an integral equation for the dynamics of a perturbation from the fixed point [47],

$$\begin{aligned}\vec{r}(\mathbf{x}, t) = &\iint_\Gamma \int_0^\infty A(\tau) W(\mathbf{u}) \vec{r}(\mathbf{x} - \mathbf{u}, t - \tau) d\tau d\mathbf{u} \\ &+ \iint_\Gamma \int_0^\infty B(\tau) U(\mathbf{u}) \vec{r}(\mathbf{x} - \mathbf{u}, t - \tau) d\tau d\mathbf{u}.\end{aligned} \tag{2}$$

as illustrated in Fig. 2a. The matrix kernels, $A(\tau)$ and $B(\tau)$, quantify excitatory and inhibitory neurons' linear response to perturbations in their input mean and variance [48]. Eq. (2) can capture an arbitrary linear dependence of firing rates on recent history, which is generally not possible in a finite system of integro-differential equations like Eq. (1). Transitioning to the temporal Laplace and spatial Fourier domains in Eq. (2) gives

$$\det\left[\widehat{A}(\lambda)\widetilde{W}(\mathbf{n}) + \widehat{B}(\lambda)\widetilde{U}(\mathbf{n}) - Id\right] = 0 \tag{3}$$

where Id is the 2×2 identity matrix, $\widetilde{W}(\mathbf{n})$ and $\widetilde{U}(\mathbf{n})$ are Fourier coefficients of $W(\mathbf{u})$ and $U(\mathbf{u})$, and $\widehat{A}(\lambda)$ and $\widehat{B}(\lambda)$ are matrices of susceptibility functions [49], which can be computed under the diffusion approximation using a Fokker-Planck formalism [20–22, 30]. Solutions, λ, to Eq. (3) are eigenvalues of the rate dynamics and the associated Fourier modes, $\mathbf{n}$, are eigenmodes.

Computation of the eigenvalues confirms that the uniform fixed point rates are stable for the network with local inhibition (Fig. 2b), but unstable for the network with broader inhibition (Fig. 2c). The eigenvalues with positive real part have non-zero imaginary part (Fig. 2c) and are associated with spatially non-uniform eigenmodes ($\mathbf{n} \neq 0$), implying a Turing-Hopf bifurcation that produces spatially coherent, time-varying patterns [17, 18]. Varying the width of inhibition shows that eigenvalues with positive real part emerge once inhibition is about twice as broad as excitation, coinciding with the emergence of high-amplitude firing rate variability in simulations (Fig. 2d, compare green and purple). Stability can also be modulated by the strength of external input to inhibitory neurons [50], showing that the dynamical state of the network can be controlled by input.

So far, we have considered purely spontaneous firing rate dynamics. We next added a time-varying external input shared by all neurons in the network (Fig. 3a). For the stable network, local firing rates approximately tracked the shared input with the addition of irregular fluctuations (Fig. 3b), consistent with Poisson-like spike-timing variability [51]. Applying principal component

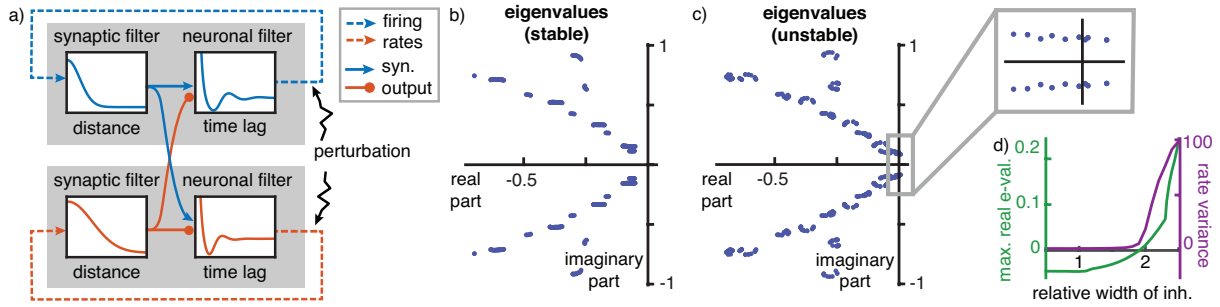

FIG. 2: **Stability of spatially extended spiking networks. a)** Schematic of stability analysis. A perturbation applied to the firing rates (dashed lines, blue for excitatory and red for inhibitory) is filtered spatially by synaptic divergence (left boxes, showing connection probability as a function of distance) to determine perturbations of synaptic currents (solid lines), which are filtered temporally by neurons (right boxes, showing linear response kernels). **b)** When inhibition is more local than excitation (as in Fig. 1a-c), all eigenvalues have negative real parts. **c)** Same as (b), except for broad inhibition (as in Fig. 1d-f). Some eigenvalues have positive real part (insert). **d)** Maximum real part of the eigenvalues (green) and the average temporal variance of the firing rate readouts (purple; units Hz^2) as a function of the relative width of inhibitory projections ($\sigma_\mathrm{i}/\sigma_\mathrm{e}$).

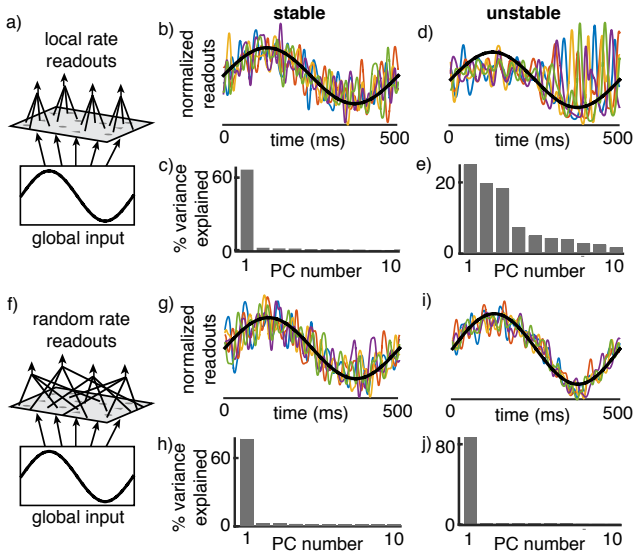

FIG. 3: **Firing rate response to global input. a)** A spatially uniform sinusoidal input was provided to all neurons in the network ($F^\mathrm{e}(t) = 3 + 1.5\sin(2\pi t)$, $F^\mathrm{i}(t) = 2.3 + 1.5\sin(2\pi t)$). **b)** Five randomly chosen firing rate readouts (colored curves) track the input (black curve) with Poisson-like variability. All curves normalized by subtracting their mean and dividing by the standard deviation. **c)** Percent variance in the 100 firing rate readouts explained by the first ten principal component projections. **d,e)** Same as b and c, but for the unstable network. **f-j)** Same as a-e except that firing rates are read out randomly and globally from the network.

(PC) analysis to the local firing rates revealed that the majority of firing rate variability is captured by the first PC projection (Fig. 3c), representing the variability inherited from the one-dimensional external input. The remaining variability was spread among higher PC projections, representing unstructured, Poisson-like spike-timing variability.

The unstable network exhibited a starkly different response to the external input. While local rates were affected by the external input, they did not reliably track it (Fig. 3d). The input evoked a high-dimensional response, with variability distributed across several PC projections (Fig. 3e). These results show that the unstable network generates high-dimensional firing rate dynamics in response to a one-dimensional input, while the stable network simply tracks the input with Poisson-like spike timing variability.

Interestingly, while *local* rate readouts of the unstable network did not track the input, random *global* readouts from the same network do track input. Specifically, we computed firing rate readouts generated from 400 excitatory neurons randomly selected from the entire network (Fig. 3f), instead of locally. These random readouts from both the stable and unstable networks reliably tracked external input (Fig. 3g-j). This finding can be understood by noting that the random readouts estimate the network-averaged rates. Eq. (3) is identical for the stable and unstable networks at the uniform eigenmode, $\mathbf{n} = \mathbf{0}$, so the networks have the same eigenvalues at that mode. Hence, the global average firing rate exhibits similar dynamics in both networks.

For the firing rate dynamics generated by the unstable network to perform reliable nonlinear computations on inputs, the response of the network should be consistent across repeated presentations of the same input. We found that the transformation of spatially uniform input considered in Fig. 3d,e was not reliable: While the first PC projection reliably tracked the input, other components were highly unreliable from trial to trial (Fig. 4a,b). We conjectured that this unreliability is due to the spatial symmetry of the network: Since the activity patterns generated by the unstable network arise through a symmetry-breaking dynamic, each time the input is presented, there are numerous firing rate patterns that

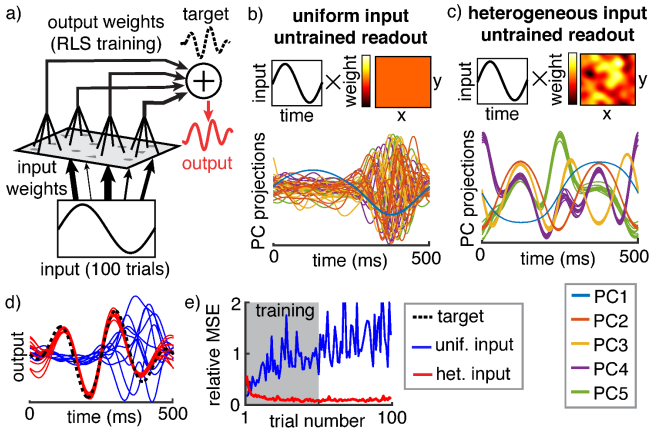

FIG. 4: **Reliable computations require heterogeneous input.** **a)** Schematic. Same as Fig. 3a except the input was repeated for 100 consecutive trials and multiplied by fixed, location-dependent weights. Readouts were multiplied by output weights that were trained to produce a target output. **b)** Overlaid plots of the first five PC projections of untrained readouts (uniform readout weights) over ten randomly selected trials when inputs were spatially uniform (same model as Fig. 3d). **c)** Same as (b) with spatially heterogeneous input weights. **d)** Trained output from last ten trials with uniform (blue) and heterogeneous (red) input weights, compared to target (dashed black). **e)** Mean-squared error of readouts.

are equally likely to be evoked. As a result, the evoked response depends on small differences in the network state when the input arrives.

We therefore considered an input that projects to the network with weights that vary across space [52] (Fig. 4a,c). This modification had a striking effect on the network response. Unlike the response to spatially uniform input, the response to spatially heterogeneous input was highly reliable from one presentation of the stimulus to another (Fig. 4c).

We next asked whether the unstable network could be trained to implement dynamical computations by using the local rate readouts as the "reservoir" in a reservoir computing model. Specifically, the local rate readouts were linearly combined to produce an output time series. The weights for the linear combination of readouts were trained using a recursive least-squares algorithm [31] that iteratively updates weights to mold the output to a target time-series [53] (Fig. 4a).

When this algorithm was applied to the firing rates produced by spatially uniform inputs (from Fig. 4b), the outputs did not produce the target time series (Fig. 4d,e, blue curves), due to the unreliability of the network response. When the same algorithm was applied to the rates produced by spatially heterogeneous inputs (from Fig. 4c), the outputs closely matched the target (Fig. 4d,e, red curves). Further simulations confirm that the network can produce a variety of target outputs from a variety of inputs [54].

Discussion There is an extensive literature on the analysis of spatially extended neural fields [23–25] and the dynamics of spiking neuron models [32, 33], but these topics are rarely combined. Previous studies found spatiotemporal dynamics in spiking networks with synaptic kinetics or delays [26, 34, 35]. Since the resonance for a Turing-Hopf bifurcation arises primarily from synaptic dynamics in these models, their stability can be captured by differential neural field equations, similar to Eq. (1). The Turing-Hopf bifurcation observed here and in previous work [9] arises from the resonance of spiking neurons, which is not captured by Eq. (1). We showed that spatial dynamics arising from the resonance of individual neurons are rendered mathematically tractable by extending linear response techniques developed for homogeneous networks [29]. This approach is applicable to the growing class of neuron models for which the linear response function can be computed [20–22].

Only a few studies have implemented reservoir computing with spiking networks. Maass et al. [36] used a spatially extended spiking network for reservoir computing, but did not explain the role of the spatial topology, which we have clarified. More recent studies [37, 38] showed that precisely tuning a sub-network of slow synapses offline can produce intricate rate dynamics in spiking networks. One of those studies [38], implemented reservoir computing with a spiking network. In the other study [37] this was only done for a rate network version of the model. It remains to be shown how this precise tuning of synapses could be achieved biologically, but inhibitory plasticity is one possibility [39].

Ostojic [40] showed that spiking networks can produce high-dimensional rate dynamics when synapses are strong, analogous to rate networks with similar structure [41], but the reliability of these dynamics and their computational capabilities were not explored. The combination of strong coupling with spatial network topology is a promising direction for future study.

Distance-dependent connectivity is ubiquitous in the brain [13–16]. We showed that this spatial topology imparts spiking neural networks with the ability to perform dynamical computations (Fig. 4) while maintaining the ability to accurately track network input (Fig. 3i,j). Hence, spatial network architecture provides a critical link between biological realism and computational capability in recurrent neural network models.

Spatially extended networks are often modeled with integro-differential equations that do not capture the history-dependence of rate dynamics. We showed that this shortcoming is overcome using linear response theory to replace the integro-differential equation, (1), with an integral equation, (2). This approach has applications in any stochastic system with spatially and temporally nonlocal interactions such as models of social networks, population dynamics and epidemiology.

We thank Ashok Litwin-Kumar, Bard Ermentrout and Brent Doiron for helpful conversations. This work was supported by NSF grant DMS-1517828.

[1] W. R. Softky and C. Koch, J Neurosci **13**, 334 (1993).
[2] M. N. Shadlen and W. T. Newsome, Curr Op Neurobiol **4**, 569 (1994).
[3] C. van Vreeswijk and H. Sompolinsky, Science **274**, 1724 (1996).
[4] C. van Vreeswijk and H. Sompolinsky, Neural Comput **10**, 1321 (1998).
[5] N. Brunel, J Comput Neurosci **8**, 183 (2000).
[6] A. Renart, J. de La Rocha, P. Bartho, L. Hollender, N. Parga, A. Reyes, and K. Harris, Science **327**, 587 (2010).
[7] M. Monteforte and F. Wolf, Phys. Rev. X **2**, 041007 (2012).
[8] G. Lajoie, K. K. Lin, and E. Shea-Brown, Phys. Rev. E **87**, 052901 (2013).
[9] R. Rosenbaum and B. Doiron, Phys Rev X **4**, 021039 (2014).
[10] R. Pyle and R. Rosenbaum, Phys. Rev. E **93**, 040302 (2016).
[11] M. M. Churchland, J. P. Cunningham, M. T. Kaufman, J. D. Foster, P. Nuyujukian, S. I. Ryu, and K. V. Shenoy, Nature **487**, 51 (2012).
[12] K. V. Shenoy, M. Sahani, and M. M. Churchland, Annu. Rev. Neurosci. **36**, 337 (2013).
[13] J. S. Lund, A. Angelucci, and P. C. Bressloff, Cereb Cortex **13**, 15 (2003).
[14] R. B. Levy and A. D. Reyes, J Neurosci **32**, 5609 (2012).
[15] M. Ercsey-Ravasz, N. T. Markov, C. Lamy, D. C. VanEssen, K. Knoblauch, Z. Toroczkai, and H. Kennedy, Neuron **80**, 184 (2013).
[16] K. Knoblauch, M. Ercsey-Ravasz, H. Kennedy, and Z. Toroczkai, in *Micro-, meso-and macro-connectomics of the brain*, edited by H. Kennedy, D. C. Van Essen, and Y. Christen (Springer, Heidelberg, 2014) pp. 45–74.
[17] M. R. Ricard and S. Mischler, J Nonlinear Sci **19**, 467 (2009).
[18] A. Roxin, N. Brunel, and D. Hansel, Phys Rev Lett **94**, 238103 (2005).
[19] D. Amit and N. Brunel, Cereb Cortex **7**, 237 (1997).
[20] M. J. E. Richardson, Phys Rev E **76**, 021919 (2007).
[21] M. J. E. Richardson, Phy Rev E **80**, 021928 (2009).
[22] R. Rosenbaum, Front Comput Neurosci **10**, 39 (2016).
[23] B. Ermentrout, Rep. Prog. Phys **61**, 353 (1998).
[24] S. Coombes, Biol Cybern **93**, 91 (2005).
[25] P. C. Bressloff, J Phys A **45**, 033001 (2012).
[26] S. Sadeh and S. Rotter, PLoS One **9**, e114237 (2014).
[27] S. Sadeh, C. Clopath, and S. Rotter, PLoS One **10**, e0127547 (2015).
[28] B. Lindner, Phys. Rep. **392**, 321 (2004).
[29] E. Ledoux and N. Brunel, Front Comput Neurosci **5**, 25 (2011).
[30] H. Risken, *The Fokker-Planck Equation* (Springer-Verlag, Berlin, 1989).
[31] D. Sussillo and L. F. Abbott, Neuron **63**, 544 (2009).
[32] W. Gerstner and W. M. Kistler, *Spiking neuron models: Single neurons, populations, plasticity* (Cambridge university press, 2002).
[33] E. M. Izhikevich, *Dynamical systems in neuroscience* (MIT press, 2007).
[34] S. Lim and M. S. Goldman, Nat. Neurosci. **16**, 1306 (2013).
[35] A. Keane and P. Gong, J Neurosci **35**, 1591 (2015).
[36] W. Maass, T. Natschläger, and H. Markram, Neural Comput **14**, 2531 (2002).
[37] G. Hennequin, T. P. Vogels, and W. Gerstner, Neuron **82**, 1394 (2014).
[38] L. F. Abbott, B. Depasquale, and R.-m. Memmesheimer, Nat. Neurosci. **19**, 350 (2016).
[39] T. P. Vogels, H. Sprekeler, F. Zenke, C. Clopath, and W. Gerstner, Science **334**, 1569 (2011).
[40] S. Ostojic, Nat. Neurosci. **17**, 594 (2014).
[41] H. Sompolinsky, A. Crisanti, and H. J. Sommers, Phys. Rev. Lett. **61**, 259 (1988).
[42] Membrane potential of neuron j in population $a = \mathrm{e}, \mathrm{i}$ obeys $V' = -(V - E_L)/\tau_m + I_j^a$ where $E_L = -70$ mV, $\tau_m = 20$ ms and each time V exceeds $V_{th} = -50$ mV, a spike is recorded and V is reset to $V_{re} = -75$ mV. In all simulations, average connection probabilities were $\overline{p}_{\mathrm{ee}} = \overline{p}_{\mathrm{ie}} = 0.0125$, $\overline{p}_{\mathrm{ei}} = \overline{p}_{\mathrm{ii}} = 0.05$, so that each neuron received 500 excitatory and 500 inhibitory inputs on average and connection weights were $j_{\mathrm{ee}} = 0.1$, $j_{\mathrm{ie}} = 0.2$ and $j_{\mathrm{ei}} = j_{\mathrm{ii}} = -0.25$ mV.
[43] $W = [w_{\mathrm{ee}}(\mathbf{u})\ w_{\mathrm{ei}}(\mathbf{u});\ w_{\mathrm{ie}}(\mathbf{u})\ w_{\mathrm{ii}}(\mathbf{u})]$ is a 2×2 matrix function where $w_{ab}(\mathbf{u}) = j_{ab} p_{ab}(\mathbf{u}) N_b$ and similarly for $U(\mathbf{u})$, except that $u_{ab} = j_{ab} w_{ab}$.
[44] See Supplementary Figure 1 for a comparison with Poisson spike trains.
[45] See Supplementary Figure 2.
[46] See Supplementary Note S.1 and Supplementary Figure 4 for a review of the stability analysis of the neural field model.
[47] See Supplementary Note S.2 for a derivation.
[48] $A(\tau) = [A_{\mathrm{e}}(\tau)\ 0;\ 0\ A_{\mathrm{i}}(\tau)]$ where $A_a(\tau)$ is the excitatory neurons' linear response to a perturbation in mean input, and similarly for $B(\tau)$.
[49] $\widetilde{W} = [\widetilde{w}_{\mathrm{ee}}\ \ \widetilde{w}_{\mathrm{ei}};\ \ \widetilde{w}_{\mathrm{ie}}\ \ \widetilde{w}_{\mathrm{ii}}]$ where $\widetilde{w}_{ab}(\mathbf{n}) = j_{ab} N_b \overline{p}_{ab} \exp(-2\pi^2 \|\mathbf{n}\|^2 \sigma_b^2)$, similarly for $\widetilde{U}$ and $\widetilde{u}_{ab}(\mathbf{n}) = j_{ab} \widetilde{w}_{ab}(\mathbf{n})$. Also, $\widehat{A}(\lambda) = [\widehat{A}_{\mathrm{e}}(\lambda)\ 0;\ 0\ \widehat{A}_{\mathrm{e}}(\lambda)]$ where $\widehat{A}_x(\lambda)$ is the susceptibility of excitatory neurons and similarly for $\widehat{B}(\lambda)$.
[50] See Supplementary Figure 3.
[51] See Supplementary Figure 1 for a comparison with Poisson spike trains.
[52] $F_j^{\mathrm{e}}(t) = 3 + 1.5 \sin(2\pi t) Q(x, y)$ where $Q(x, y)$ is a realization of quenched, spatial Gaussian noise with unit variance and $E[Q(\mathbf{x})Q(\mathbf{x} + \mathbf{u})] = \exp(-|\mathbf{u}|^2/(2\sigma^2))$ where $\sigma = 0.1$.
[53] See Supplementary Note S.3 for description of numerical procedure for training weights.
[54] See Supplementary Figure 5 for further reservoir computing examples